JMLR: Workshop and Conference Proceedings vol 65:1–16, 2017

# On the effect of pooling on the geometry of representations

**Gary Bécigneul** GARY.BECIGNEUL@INF.ETHZ.CH
*Department of Computer Science, ETH Zürich and Max Planck ETH Center for Learning Systems*

arXiv:1703.06726v1 [cs.LG] 20 Mar 2017

**Editor:** Under Review for COLT 2017

## Abstract

In machine learning and neuroscience, certain computational structures and algorithms are known to yield disentangled representations without us understanding why, the most striking examples being perhaps convolutional neural networks and the ventral stream of the visual cortex in humans and primates. As for the latter, it was conjectured that representations may be disentangled by being flattened progressively and at a local scale (DiCarlo and Cox, 2007). An attempt at a formalization of the role of invariance in learning representations was made recently, being referred to as *I*-theory (Anselmi et al., 2013b). In this framework and using the language of differential geometry, we show that pooling over a group of transformations of the input contracts the metric and reduces its curvature, and provide quantitative bounds, in the aim of moving towards a theoretical understanding on how to disentangle representations.

**Keywords:** Differential Geometry, *I*-theory, deep learning, pooling, disentangle, representation, curvature, group orbit

## 1. Introduction

What does disentangling representations mean? In machine learning and neurosciences, representations being tangled has two principal interpretations, and they are intimately connected with each other. The first one is geometrical: consider two sheets of paper of different colors, place one of the two on top of the other, and crumple them together in a paper ball; now, it may look difficult to separate the two sheets with a third one: they are *tangled*, one color sheet representing one class of a classification problem. The second one is analytical: consider a dataset being parametrized by a set of coordinates $\{x_i\}_{i\in I}$, such as images parametrized by pixels, and a classification task between two classes of images. On the one hand, we cannot find a subset $\{x_i\}_{i\in J}$ with $J \subset I$ of this coordinate system such that a variation of these would not change the class of an element, while still spanning a reasonable amount of different images of this class. On the other hand, we are likely to be capable of finding a large amount of transformations preserving the class of any image of the dataset, without being expressible as linear transformations on this coordinate system, and this is another way to interpret representations or factors of variation as being tangled.

Why is disentangling representations important? On the physiological side, the brains of humans and primates alike have been observed to solve object recognition tasks by progressively disentangling their representations *via* the visual stream, from V1 to the IT cortex (DiCarlo and Cox, 2007; DiCarlo et al., 2012). On the side of deep learning, deep convolutional neural networks are also able to disentangle highly tangled representations, since a softmax − which, geometrically, performs essentially a linear separation − computed on the representation of their last hidden layer

© 2017 G. Bécigneul.

can yield very good accuracy (Krizhevsky et al., 2012). Conversely, disentangling representations might be sufficient to pre-solve practically any task relevant to the observed data (Bengio, 2013).

How can we design algorithms in order to move towards more disentangled representations? Although it was conjectured that the visual stream might disentangle representations by flattening them locally, thus inducing a decrease in the curvature globally (DiCarlo and Cox, 2007), the mechanisms underlying such a disentanglement, whether it be for the brain or deep learning architectures, remain very poorly understood (DiCarlo and Cox, 2007; Bengio, 2013). However, it is now of common belief that computing representations that are invariant with respect to irrelevant transformations of the input data can help. Indeed, on the one hand, deep convolutional networks have been noticed to naturally learn more invariant features with deeper layers (Goodfellow et al., 2009; Lenc and Vedaldi, 2015; Tensmeyer and Martinez, 2016). On the other hand, the V1 part of the brain similarly achieves invariance to translations and rotations *via* a "pinwheels" structure, which can be seen as a principal fiber bundle (Petitot, 2003; Poggio et al., 2012). Conversely, enforcing a higher degree of invariance with respect to not only translations, but also rotations, flips, and other groups of transformation has been shown to achieve state-of-the-art results in various machine learning tasks (Bruna and Mallat, 2013; Gens and Domingos, 2014; Oyallon and Mallat, 2015; Dieleman et al., 2016; Cohen and Welling, 2016a,b), and is believed to help in linearizing small diffeomorphisms (Mallat, 2016). To the best of our knowledge, the main theoretical efforts in this direction include the theory of scattering operators (Mallat, 2012; Wiatowski and Bölcskei, 2015) as well as $I$-theory (Anselmi et al., 2013b,a; Anselmi and Poggio, 2014; Anselmi et al., 2016). In particular, $I$-theory permits to use the whole apparatus of kernel theory to build invariant features (Mroueh et al., 2015; Raj et al., 2016).

Our work builds a bridge between the idea that disentangling is a result of *(i)* a local decrease in the curvature of the representations, and *(ii)* building representations that are invariant to nuisance deformations, by proving that pooling over such groups of transformations results in a local decreasing of the curvature.

We start by providing some background material, after which we introduce our formal framework and theorems, which we then discuss in the case of the non-commutative group generated by translations and rotations.

## 2. Some background material

### 2.1. Groups and geometry

A *group* is a set $G$ together with a map $\cdot : G \times G \to G$ such that:
*(i)* $\forall g, g', g'' \in G,\ g \cdot (g' \cdot g'') = (g \cdot g') \cdot g''$,
*(ii)* $\exists e \in G,\ \forall g \in G,\ g \cdot e = e \cdot g = g$,
*(iii)* $\forall g \in G,\ \exists g^{-1} \in G :\ g \cdot g^{-1} = g^{-1} \cdot g = e$,
where $e$ is called the *identity element*. We write $gg'$ instead of $g \cdot g'$ for simplicity. If, moreover, $gg' = g'g$ for all $g, g' \in G$, then $G$ is said to be *commutative* or *abelian*.

A *subgroup* of $G$ is a set $H \subset G$ such that for all $h, h' \in H$, $hh' \in H$ and $h^{-1} \in H$. A subgroup $H$ of a group $G$ is said to be *normal* in $G$ if for all $g \in G$, $gH = Hg$, or equivalently, for all $g \in G$ and $h \in H$, $ghg^{-1} \in H$. If $G$ is abelian, then all of its subgroups are normal in $G$.

A *Lie group* is a group which is also a smooth manifold, and such that its product law and inverse map are smooth with respect to its manifold structure. A Lie group is said to be *locally compact* if each of its element possesses a compact neighborhood. On every locally compact Lie group, one can define a *Haar measure*, which is a left-invariant, non-trivial Lebesgue measure on its Borel algebra, and is uniquely defined up to a positive scaling constant. If this Haar measure is also right-invariant, then the group is said to be *unimodular*. This Haar measure is always finite on compact sets, and strictly positive on non-empty open sets. Examples of unimodular Lie groups include in particular all abelian groups, compact groups, semi-simple Lie groups and connected nilpotent Lie groups.

A group $G$ is said to be *acting* on a set $X$ if we have a map $\cdot : G \times X \to X$ such that for all $g, g' \in G$, for all $x \in X$, $g \cdot (g' \cdot x) = (gg') \cdot x$ and $e \cdot x = x$. If this map is also smooth, then we say that $G$ is *smoothly acting* on $X$. We write $gx$ instead of $g \cdot x$ for simplicity. Then, the *group orbit* of $x \in X$ under the action of $G$ is defined by $G \cdot x = \{gx \mid g \in G\}$, and the stabilizer of $x$ by $G_x = \{g \in G \mid gx = x\}$. Note that $G_x$ is always a subgroup of $G$, and that for all $x, y \in X$, we have either $(G \cdot x) \cap (G \cdot y) = \varnothing$, or $G \cdot x = G \cdot y$. Hence, we can write $X$ as the disjoint union of its group orbits, *i.e.* there exists a minimal subset $\tilde{X} \subset X$ such that $X = \sqcup_{x \in \tilde{X}} G \cdot x$. The set of orbits of $X$ under the action of $G$ is written $X/G$, and is in one-to-one correspondence with $\tilde{X}$. Moreover, note that if $H$ is a subgroup of $G$, then $H$ is naturally acting on $G$ via $(h, g) \in H \times G \mapsto hg \in G$; if we further assume that $H$ is normal in $G$, then one can define a canonical group structure on $G/H$, thus turning the canonical projection $g \in G \mapsto H \cdot g$ into a group morphism.

A *diffeomorphism* between two manifolds is a map that is smooth, bijective and has a smooth inverse. A *group morphism* between two groups $G$ and $G'$ is a map $\varphi : G \to G'$ such that for all $g_1, g_2 \in G$, $\varphi(g_1 g_2) = \varphi(g_1)\varphi(g_2)$. A *group isomorphism* is a bijective group morphism, and a *Lie group isomorphism* is a group isomorphism that is also a diffeomorphism.

The *Lie algebra* $\mathfrak{g}$ of a Lie group $G$ is its tangent space at $e$, and is endorsed with a bilinear map $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$ called its *Lie bracket*, and such that for all $x, y, z \in \mathfrak{g}$, $[x, y] = -[y, x]$ and $[x, [y, z]] + [y, [z, x]] + [z, [x, y]] = 0$. Moreover, there is a bijection between $\mathfrak{g}$ and left-invariant vector fields on $G$, defined by $\xi \in \mathfrak{g} \mapsto \{g \in G \mapsto d_e L_g(\xi)\}$, where $L_g(h) = gh$ is the left translation. Finally, the flow $t \mapsto \phi_t$ of such a left-invariant vector field $X_\xi$ is given by $\phi_t(g) = g \exp(t\xi)$, where $\exp : \mathfrak{g} \to G$ is the exponential map on $G$.

For more on Lie groups, Lie algebras, Lie brackets and group representations, see Kirillov (2008), and for a rapid and clear presentation of the notions of sectional curvature and Riemannian curvature, see Andrews and Hopper (2010).

### 2.2. *I*-theory

*I*-theory aims at understanding how to compute a representation of an image $I$ that is both unique and invariant under some deformations of a group $G$, and how to build such representations in a hierarchical way (Poggio et al., 2012; Anselmi et al., 2013b,a; Anselmi and Poggio, 2014; Anselmi et al., 2016).

Suppose that we are given a Hilbert space $\mathcal{X}$, typically $L^2(\mathbb{R}^2)$, representing the space of images. Let $G$ be a locally compact group acting on $\mathcal{X}$. Then, note that the group orbit $G \cdot I$ constitutes such an invariant and unique representation of $I$, as $G \cdot I = G \cdot (gI)$, for all $g \in G$, and since two group orbits intersecting each other are equal.

But how can we compare such group orbits? For an image $I \in \mathcal{X}$, define the map $\Theta_I : g \in G \mapsto gI \in \mathcal{X}$ and the probability distribution $P_I(A) = \mu_G(\Theta_I^{-1}(A))$ for any borel set $A$ of $\mathcal{X}$, where $\mu_G$ is the Haar measure on $G$. For $I, I' \in \mathcal{X}$, write $I \sim I'$ is there exists $g \in G$ such that $I = gI'$. Then, one can prove that $I \sim I'$ if and only if $P_I = P_{I'}$. Hence, we could compare $G \cdot I$ and $G \cdot I'$ by comparing $P_I$ and $P_{I'}$. However, computing $P_I$ can be difficult, so one must be looking for ways to approximate $P_I$. If $t \in \mathbb{S}(L^2(\mathbb{R}^2))$, define $P_{\langle I,t \rangle}$ to be the distribution associated with the random variable $g \mapsto \langle gI, t \rangle$. One can then prove that $P_I = P_{I'}$ if and only if $P_{\langle I,t \rangle} = P_{\langle I',t \rangle}$ for all $t \in \mathbb{S}(L^2(\mathbb{R}^2))$, and then provide a lower bound on the sufficient number $K$ of such templates $t_k$, $1 \leqslant k \leqslant K$, drawn uniformly on $\mathbb{S}(L^2(\mathbb{R}^2))$, in order to recover the information of $P_I$ up to some error $\varepsilon$ and with high probability $1 - \delta$. Finally, each $P_{\langle I,t_k \rangle}$ can be approximated by a histogram

$$h_n^k(I) = \frac{1}{|G|} \sum_{g \in G} \eta_n(\langle gI, t_k \rangle),$$

if $G$ is finite or

$$h_n^k(I) = \frac{1}{\mu_G(G)} \int_{g \in G} \eta_n(\langle gI, t_k \rangle) d\mu_G(g),$$

if $G$ is compact, where $\eta_n$ are various non-negative and possibly non-linear functions, $1 \leqslant n \leqslant N$, such as sigmoid, ReLU, modulus, hyperbolic tangent or $x \mapsto |x|^p$, among others.

In the case where the group $G$ is only partially observable (for instance if $G$ is only locally compact but not bounded), one can define instead a "partially invariant representation", replacing each $h_n^k(I)$ by

$$\frac{1}{\mu_G(G_0)} \int_{g \in G_0} \eta_n(\langle gI, t_k \rangle) d\mu_G(g),$$

where $G_0$ is a compact subset of $G$ which can be observed in practice. Under some "localization condition" (see (Anselmi et al., 2013b)), it can be proved that this representation is invariant under deformations by elements of $G_0$. When this localization condition is not met, we do not have any exact invariance a priori, but one might expect that the variation in directions defined by the symmetries of $G_0$ is going to be reduced.

For instance, let $G$ be the group $\mathbb{R}^2$ of translations in the plane, $G_0 = [-a, a]^2$ for some $a > 0$, $\eta : x \mapsto (\sigma(x))^2$ where $\sigma$ is a point-wise non-linearity commonly used in neural networks, and $t_k \in \mathbb{S}(L^2(\mathbb{R}^2))$ for $1 \leqslant k \leqslant K$. Then, note that the quantities

$$\sqrt{|G_0| h_k(I)} = \sqrt{\sum_{g \in G_0} \eta(\langle gI, t_k \rangle)},$$

for $1 \leqslant k \leqslant K$ are actually computed by a 1-layer convolutional neural network with filters $(t_k)_{1 \leqslant k \leqslant K}$, non-linearity $\sigma$ and $L^2$-pooling. Moreover, the family $(\sqrt{|G_0| h_k(gI)})_{g \in G}$ is exactly

the output of this convolutional layer, thus describing a direct correspondence between pooling and locally averaging over a group of transformations.

Another correspondence can be made between this framework and deep learning architectures. Indeed, assume that during learning, the set of filters of a layer of a convolutional neural network becomes stable under the action of some unknown group $G$ acting on the pixel space, and denote by $\sigma$ the point-wise non-linearity computed by the network. Moreover, suppose that the convolutional layer and point-wise non-linearity are followed by an $L^p$-pooling, defined by $\Pi_\phi^p(I)(x) = \left(\int_{y\in\mathbb{R}^2} |I(y)\mathbf{1}_{[0,a]^2}(x-y)|^p \, dy\right)^{1/p}$. Then, observe that the convolutional layer outputs the following feature maps:

$$\{\Pi_\phi^p(\sigma(I \star t_k))\}_{1\leqslant k\leqslant K}.$$

Besides, if the group $G$ has a unitary representation, and if its action preserves $\mathbb{R}^2$, then for all $g \in G$ and $1 \leqslant k \leqslant K$, we have

$$\Pi_\phi^p(\sigma(gI \star t_k)) = \Pi_\phi^p(\sigma(g(I \star g^{-1}t_k))) = g\Pi_\phi^p(\sigma(I \star g^{-1}t_k)).$$

Then, the following layer of the convolutional network is going to compute the sum across channels $k$ of these different signals. However, if our set of filters $t_k$ can be written as $G_0 \cdot t$ for some filter $t$ and a subpart $G_0$ of $G$, then this sum will be closely related to a histogram as in *I*-theory:

$$\sum_{g\in G_0} \Pi_\phi^p(\sigma(I \star gt)) = \sum_{g\in G_0} g\Pi_\phi^p(\sigma(g^{-1}I \star t_k)).$$

In other words, (local) group invariances are free to appear during learning among filters of a convolutional neural network, and will naturally be pooled over by the next layer. For more on this, see (Bruna et al., 2013; Mallat, 2016).

Finally, let's mention that this implicit pooling over symmetries can also be computed explicitly, and such group invariances across filters enforced, if we know the group in advance, as in G-CNNs and steerable CNNs (Cohen and Welling, 2016a,b).

## 3. Main results: formal framework and theorems

Let $G$ be a finite-dimensional, locally compact and unimodular Lie group smoothly acting on $\mathbb{R}^2$. This defines an action $(L_g f)(x) = f(g^{-1}x)$ on $L^2(\mathbb{R}^2)$. Let $G_0$ be a compact neighborhood of the identity element $e$ in $G$, and assume that there exists $\lambda > 0$ such that for all $g_0 \in G_0$, $\sup_{x\in\mathbb{R}^2} |J_{g_0}(x)| \leqslant \lambda$, where $J_g$ is the determinant of the Jacobian matrix of $g$ seen as a diffeomorphism of $\mathbb{R}^2$. We define $\Phi : L^2(\mathbb{R}^2) \to L^2(\mathbb{R}^2)$, the averaging operator on $G_0$, by

$$\Phi(f) = \frac{1}{\mu_G(G_0)} \int_{g\in G_0} L_g f \, d\mu_G(g).$$

Our first result describes how the euclidean distance in $L^2(\mathbb{R}^2)$ between a function $f$ and its translation by some $g \in G_0$ is contracted by this locally averaging operator.

**Theorem 1.**
For all $f \in L^2(\mathbb{R}^2)$, for all $g \in G$,

$$\|\Phi(L_g f) - \Phi(f)\|_2 \leqslant \sqrt{\lambda} \max\left(1, \sqrt{\|J_g\|_\infty}\right) \frac{\mu_G((G_0 g)\Delta G_0)}{\mu_G(G_0)} \|f\|_2.$$

*Proof:* See Appendix A.

The symbol $\Delta$ above is defined $A\Delta B = (A \cup B) \setminus (A \cap B) = (A \setminus B) \cup (B \setminus A)$. Note that, as one could have expected, this result doesn't depend on the scaling constant of the Haar measure. Intuitively, this result formalizes the idea that locally averaging with respect to some factors of variation, or coordinates, will reduce the variation with respect to those coordinates. The following drawings illustrate the intuition behind Theorem 1, where we pass from left to right by applying $\Phi$.

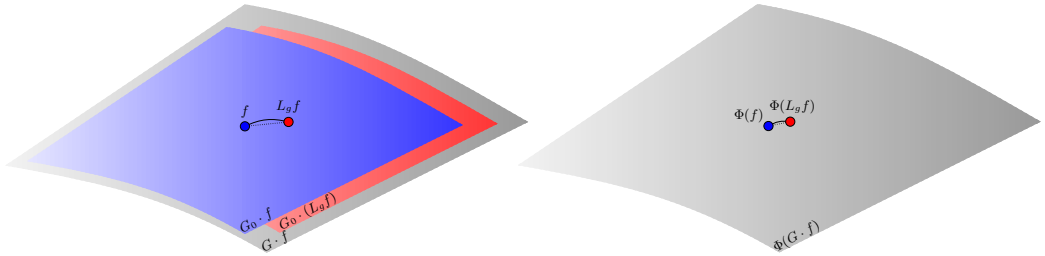


Figure 1: Concerning the drawing on the left-hand side, the blue and red areas correspond to the compact neighborhood $G_0$ centered in $f$ and $L_g f$ respectively, the grey area represents only a visible subpart of the whole group orbit, the thick, curved line is a geodesic between $f$ and $L_g f$ inside the orbit $G \cdot f$, and the dotted line represents the line segment between $f$ and $L_g f$ in $L^2(\mathbb{R}^2)$, whose size is given by the euclidean distance $\|L_g f - f\|_2$.

Note that the quantity $\frac{\mu_G((G_0 g)\Delta G_0)}{\mu_G(G_0)}$, depending on the geometry of the group, is likely to decrease when we increase the size of $G_0$: if $G = \mathbb{R}^2$ is the translation group, $G_0 = [0, a]^2$ for some $a > 0$, and $g_\varepsilon$ is the translation by the vector $(\varepsilon, \varepsilon)$, then $\mu_G$ is just the usual Lebesgue measure in $\mathbb{R}^2$ and

$$\frac{\mu_G((G_0 g_\varepsilon)\Delta G_0)}{\mu_G(G_0)} \underset{\varepsilon \to 0}{\sim} 2\frac{2a\varepsilon}{a^2} = \frac{4\varepsilon}{\sqrt{\mu_G(G_0)}}.$$

Indeed, locally averaging over a wider area will decrease the variation even more.

As images are handily represented by functions from the space of pixels $\mathbb{R}^2$ to either $\mathbb{R}$ or $\mathbb{C}$, let us define our dataset $\mathcal{X}$ to be a finite-dimensional manifold embedded in a bigger space of functions $\mathcal{Y}$. As for technical reasons we will need our functions to be $L^2$, smooth, and with a gradient having a fast decay at infinity, we choose $\mathcal{Y}$ to be the set of functions $f \in L^2(\mathbb{R}^2) \cap \mathcal{C}^\infty(\mathbb{R}^2)$ such that $|\langle \nabla f(x), x \rangle| = O_{x \to \infty}(\frac{1}{\|x\|^{1+\varepsilon}})$, for some fixed small $\varepsilon > 0$. Note that in practice, images are only non-zero on a compact domain, therefore these assumptions are not restrictive.

Further assume that for all $f \in \mathcal{X}$, for all $g \in G$, $L_g f \in \mathcal{X}$. Intuitively, $\mathcal{X}$ is our manifold of images, and $G$ corresponds to the group of transformations that are not relevant to the task at hand. Recall that from $I$-theory, the orbit of an image $f$ under $G$ constitutes a good unique and invariant

representation. Here, we are interested in comparing $G \cdot f$ and $\Phi(G \cdot f)$, *i.e.* before and after locally averaging.

But how can we compute a bound on the curvature of $\Phi(G \cdot f)$? It is well known that in a Lie group endorsed with a bi-invariant pseudo-Riemannian metric $\langle \cdot, \cdot \rangle$, the Riemann curvature tensor is given by

$$R(X, Y, Z, W) = -\frac{1}{4}\langle [X, Y], [Z, W] \rangle,$$

where $X, Y, Z, W$ are left-invariant vector-fields, and hence if $(X, Y)$ forms an orthonormal basis of the plane they span, then the sectional curvature is given by

$$\kappa(X \wedge Y) = R(X, Y, Y, X) = \frac{1}{4}\langle [X, Y], [X, Y] \rangle.$$

Therefore, would we be able to define a Lie group structure and a bi-invariant pseudo-Riemannian metric on $\Phi(G \cdot f)$, we could use this formula to compute its curvature. First, we are going to define a Lie group structure on $G \cdot f$, which we will then transport on $\Phi(G \cdot f)$. As a Lie group structure is made of a smooth manifold structure and a compatible group structure, we need to construct both. In order to obtain the group structure on the orbit, let's assume that the stabilizer $G_f$ is normal; a condition that is met for instance if $G$ is abelian, or if this subgroup is trivial, meaning that $f$ does not have internal symmetries corresponding to those of $G$, which is only a technical condition, as it can be enforced in practice by slightly deforming $f$, by breaking the relevant symmetries with a small noise. Besides, in order to obtain a smooth manifold structure on the orbits, we need to assume that $G_f$ is an embedded Lie subgroup of $G$, which, from proposition B.0 (see appendix), is met automatically when this group admits a finite-dimensional representation.

Then, from proposition B.1, there is one and only one manifold structure on the topological quotient space $G/G_f$ turning the canonical projection $\pi : G \to G/G_f$ into a smooth submersion; moreover, the action of $G$ on $G/G_f$ is smooth, $G/G_f$ is a Lie group, $\pi$ is a Lie group morphism, the Lie algebra $\mathfrak{g}_f$ of $G_f$ is an ideal of the Lie algebra $\mathfrak{g}$ of $G$ and the linear map from $T_eG/T_eG_f$ to $T_{eG_f}(G/G_f)$ induced by $T_e\pi$ is a Lie algebra isomorphism from $\mathfrak{g}/\mathfrak{g}_f$ to the Lie algebra of $G/G_f$.

Finally, we need a geometrical assumption on the orbits, insuring that $G$ is warped on $G \cdot f$ in a way that is not "fractal", *i.e.* that this orbit can be given a smooth manifold structure: assume that $G \cdot f$ is locally closed in $\mathcal{X}$. Using this assumption and proposition B.2, the canonical map $\Theta_f : G/G_f \to \mathcal{X}$ defined by $\Theta_f(gG_f) = L_g f$ is a one-to-one immersion, whose image is the orbit $G \cdot f$, which is a submanifold of $\mathcal{X}$; moreover, $\Theta_f$ is a diffeomorphism from $G/G_f$ to $G \cdot f$. Further notice that $\Theta_f$ is $G$-equivariant, *i.e.* for all $g, g' \in G$,

$$\Theta_f(g(g'G_f)) = Lgg'f = L_g L_{g'} f = L_g \Theta_f(g'G_f).$$

Moreover, we can define on $G \cdot f$ a group law by

$$(L_{g_1} f) \cdot (L_{g_2} f) := L_{g_1 g_2} f,$$

for $g_1, g_2 \in G$. Indeed, let's prove that this definition doesn't depend on the choice of $g_1, g_2$. Assume that $g_i = a_i b_i$ for $a_i \in G$ and $b_i \in G_f$, $i \in \{1, 2\}$. Then, as $G_f$ is normal in $G$, there exists

$b_1' \in G_f$ such that $b_1 a_2 = a_2 b_1'$. Then $g_1 g_2 = a_1 a_2 b_1' b_2$ and hence $L_{g_1 g_2} f = L_{a_1 a_2} f$, and this group law is well-defined. Now that $G \cdot f$ is a group, observe that $\Theta_f$ is a group isomorphism from $G/G_f$ to $G \cdot f$. Indeed, it is bijective since it is a diffeomorphism, and it is a group morphism as

$$\Theta_f((gG_f)(g'G_f)) = \Theta_f((gg')G_f) = L_{gg'} f = (L_g f) \cdot (L_{g'} f) = \Theta_f(gG_f) \cdot \Theta_f(g'G_f).$$

Hence, $G \cdot f$ is also a Lie group, since $G/G_f$ is a Lie group and $\Theta_f : G/G_f \to G \cdot f$ is a diffeomorphism. Moreover, $Lie(G \cdot f)$ is isomorphic to $\mathfrak{g}/\mathfrak{g}_f$ as a Lie algebra, since they are isomorphic as vector spaces ($\Theta_f$ being an immersion), and by the fact that the pushforward of a diffeomorphism always preserves the Lie bracket.

Now that we have defined a Lie group structure on $G \cdot f$, how can we obtain one on $\Phi(G \cdot f)$? Suppose that $\Phi$ is injective on $G \cdot f$ and on $Lie(G \cdot f)$. We can thus define a group law on $\Phi(G \cdot f)$ by:

$$\forall g_1, g_2 \in G/G_f,\ \Phi(L_{g_1} f) \cdot \Phi(L_{g_2} f) := \Phi(L_{g_1 g_2} f).$$

As the inverse function theorem tells us that $\Phi$ is a diffeomorphism from $G \cdot f$ onto its image, $\Phi(G \cdot f)$ is now endorsed with a Lie group structure. However, in order to carry out the relevant calculations, we still need to define left-invariant vector-fields on our Lie group orbits.

For all $\xi \in \mathfrak{g}$, define the following left-invariant vector-fields respectively on $G \cdot f$ and $\Phi(G \cdot f)$:

$$X_\xi : L_g f \mapsto \frac{d}{dt}_{|t=0} (L_g L_{\exp(t\xi)} f),$$

$$\tilde{X}_\xi : \Phi(L_g f) \mapsto \frac{d}{dt}_{|t=0} \Phi(L_g L_{\exp(t\xi)} f).$$

We can now state the following theorem:

**Theorem 2.**
For all $f \in \mathcal{X}$, for all $\xi, \xi' \in \mathfrak{g}$,

$$\|[\tilde{X}_\xi, \tilde{X}_{\xi'}]_{\Phi(f)}\|_2^2 \leqslant \lambda \Big[\frac{d}{ds}_{|s=0} \frac{\mu_G((G_0 \exp(s[\xi, \xi']))\Delta G_0)}{\mu_G(G_0)}\Big]^2 \|f\|_2^2.$$

*Proof:* See Appendix A.

As $\mathcal{X}$ is a manifold embedded in $L^2(\mathbb{R}^2)$, it inherits a Riemannian metric by projection of the usual inner-product of $L^2(\mathbb{R}^2)$ on the tangent bundle of $\mathcal{X}$. Moreover, if we further assume that for all $g \in G$, $|J_g| = 1$, then this Riemannian metric is bi-invariant, and we can finally use the above formula on the Riemannian curvature, together with the previous inequality, to compute a bound on the curvature in a Lie group endorsed with an bi-invariant metric:

**Corollary.**
For all $f \in \mathcal{X}$, for all $\xi, \xi' \in \mathfrak{g}$,

$$0 \leqslant R_{\Phi(f)}(\tilde{X}_\xi, \tilde{X}_{\xi'}, \tilde{X}_{\xi'}, \tilde{X}_\xi) \leqslant \Big[\frac{1}{2} \frac{d}{ds}_{|s=0} \frac{\mu_G((G_0 \exp(s[\xi, \xi']))\Delta G_0)}{\mu_G(G_0)}\Big]^2 \|f\|_2^2.$$

And if $(\tilde{X}_\xi, \tilde{X}_{\xi'})$ forms an orthonormal basis of the plane they span in $Lie(\Phi(G \cdot f)) = \Phi(Lie(G \cdot f))$, then:

$$0 \leqslant \kappa_{\Phi(f)}(\tilde{X}_\xi \wedge \tilde{X}_{\xi'}) \leqslant \big[\frac{1}{2}\frac{d}{ds}_{|s=0} \frac{\mu_G((G_0 \exp(s[\xi, \xi']))\Delta G_0)}{\mu_G(G_0)}\big]^2 \|f\|_2^2.$$

*Remark.* The sectional curvature of the basis $(\tilde{X}_\xi, \tilde{X}_{\xi'})$ at $\Phi(f)$ is also the Gaussian curvature of the two-dimensional surface swept out by small geodesics induced by linear combinations of $\tilde{X}_\xi(\Phi(f))$ and $\tilde{X}_{\xi'}(\Phi(f))$.

Among well-known finite-dimensional, locally compact and unimodular Lie group smoothly acting on $\mathbb{R}^2$, there are the group $\mathbb{R}^2$ of translations, the compact groups $O(2)$ and $SO(2)$, the euclidean group $E(2)$, as well as transvections, or shears. Moreover, another class of suitable unimodular Lie groups is given by the one-dimensional flows of Hamiltonian systems, which, as deformations of images, could be interpreted as the smooth evolutions of the screen in a video over time, provided that these evolutions can be expressed as group actions on the pixel space.

Finally, let's see what Theorem 2 gives us in the case $G = \mathbb{R}^2 \ltimes SO(2)$. Note that this group is not commutative, and its curvature form is not identically zero. Let $\theta \in (-\pi, \pi)$, $a > 0$, and $G_0 = [-\theta, \theta] \times [0, a]^2$. A representation of this group is given by matrices of the form

$$g(\theta, x, y) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) & x \\ \sin(\theta) & \cos(\theta) & y \\ 0 & 0 & 1 \end{pmatrix},$$

and a representation of its Lie algebra is given by

$$\xi(\zeta, x, y) = \begin{pmatrix} 0 & -\zeta & x \\ \zeta & 0 & y \\ 0 & 0 & 0 \end{pmatrix}.$$

The Lie bracket is then given by

$$[\xi(\zeta, x, y), \xi(\zeta', x', y')] = \xi(\zeta, x, y)\xi(\zeta', x', y') - \xi(\zeta', x', y')\xi(\zeta, x, y) = \xi(0, \zeta' y - \zeta y', \zeta x' - \zeta' x).$$

As the exponential map on the group of translations is the identity map, and as the Haar measure on $\mathbb{R}^2 \ltimes SO(2)$ is just the product of the Haar measures on $\mathbb{R}^2$ and $SO(2)$, we have

$$\begin{aligned}\mu_G((G_0 \exp(s[\xi(\zeta, x, y), \xi(\zeta', x', y')]))\Delta G_0) = \\ 2\theta \mu_{\mathbb{R}^2}(([s(\zeta' y - \zeta y'), s(\zeta' y - \zeta y') + a] \times [s(\zeta x' - \zeta' x), s(\zeta x' - \zeta' x) + a])\Delta [0, a]^2),\end{aligned}$$

and $\mu_G(G_0) = 2\theta a^2$. Therefore, when $s \to 0$, we have

$$\begin{aligned}\mu_G((G_0 \exp(s[\xi(\zeta, x, y), \xi(\zeta', x', y')]))\Delta G_0) \sim \\ 2\theta \times 2(as(\zeta' y - \zeta y') + as(\zeta x' - \zeta' x)) = 4\theta as(\zeta(x' - y') - \zeta'(x - y)),\end{aligned}$$

from what we deduce that

$$\Big[\frac{1}{2}\frac{d}{ds}_{|s=0}\frac{\mu_G((G_0\exp(s[\xi(\zeta,x,y),\xi(\zeta',x',y']))\Delta G_0)}{\mu_G(G_0)}\Big]^2=\frac{(\zeta(x'-y')-\zeta'(x-y))^2}{a^2}.$$

As a consequence, if $f\in\mathcal{X}$ is an image in our dataset, of $L^2$-norm equal to 1, and if we choose $\xi(\zeta,x,y)$ and $\xi(\zeta',x',y')$ such that the $L^2$ functions $\tilde{X}_{\xi(\zeta,x,y)}(\Phi(f))$ and $\tilde{X}_{\xi(\zeta',x',y')}(\Phi(f))$ are orthogonal in $L^2$ and have $L^2$-norm equal to 1, then the Gaussian curvature $\kappa$ of the 2-dimensional surface swept out by these two vector fields around $\Phi(f)$, in the Lie group $\Phi(G\cdot f)$, is smaller than:

$$\kappa\leqslant\frac{(\zeta(x'-y')-\zeta'(x-y))^2}{a^2}.$$

## 4. Conclusion

Being able to disentangle highly tangled representations is a very important and challenging problem in machine learning. In deep learning in particular, there exist successful algorithms that may disentangle highly tangled representations in some situtations, without us understanding why. Similarly, the ventral stream of the visual cortex in humans and primates seems to perform such a disentanglement of representations, but, again, the reasons behind this process are difficult to understand. It is believed that making representations invariant to some nuisance deformations, as well as locally flattening them, might help or even be an essential part of the disentangling process. As shown by our theorems, there is a connection between these two intuitions, in the sense that achieving a higher degree of invariance with respect to some group transformations will flatten the representations in directions of the tangent space corresponding to the Lie algebra generators of these transformations. Using our theorems, we showed that in the case of the group of positive affine isometries, a precise bound on the sectional curvature can be computed, with respect to the pooling parameters. We hope that this work will encourage the geometrical study of how representations evolve during learning, in function of the hyperparameters of the algorithm that is used on these representations.

## Acknowledgments

We thank Simon Janin and Victor Godet for interesting conversations.

## References

Ben Andrews and Christopher Hopper. *The Ricci flow in Riemannian geometry: a complete proof of the differentiable 1/4-pinching sphere theorem*. Springer, 2010.

Fabio Anselmi and Tomaso Poggio. Representation learning in sensory cortex: a theory. Technical report, Center for Brains, Minds and Machines (CBMM), 2014.

Fabio Anselmi, Joel Z Leibo, Lorenzo Rosasco, Jim Mutch, Andrea Tacchetti, and Tomaso Poggio. Magic materials: a theory of deep hierarchical architectures for learning sensory representations. *CBCL paper*, 2013a.

Fabio Anselmi, Joel Z Leibo, Lorenzo Rosasco, Jim Mutch, Andrea Tacchetti, and Tomaso Poggio. Unsupervised learning of invariant representations in hierarchical architectures. *arXiv preprint arXiv:1311.4158*, 2013b.

Fabio Anselmi, Lorenzo Rosasco, and Tomaso Poggio. On invariance and selectivity in representation learning. *Information and Inference*, 5(2):134–158, 2016.

Mathieu Aubry and Bryan C Russell. Understanding deep features with computer-generated imagery. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2875–2883, 2015.

Yoshua Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.

Yoshua Bengio. Deep learning of representations: Looking forward. In *International Conference on Statistical Language and Speech Processing*, pages 1–37. Springer, 2013.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

Joan Bruna and Stéphane Mallat. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013.

Joan Bruna, Arthur Szlam, and Yann LeCun. Learning stable group invariant representations with convolutional networks. *arXiv preprint arXiv:1301.3537*, 2013.

Taco S Cohen and Max Welling. Transformation properties of learned visual representations. *arXiv preprint arXiv:1412.7659*, 2014.

Taco S Cohen and Max Welling. Group equivariant convolutional networks. *arXiv preprint arXiv:1602.07576*, 2016a.

Taco S Cohen and Max Welling. Steerable cnns. *arXiv preprint arXiv:1612.08498*, 2016b.

James J DiCarlo and David D Cox. Untangling invariant object recognition. *Trends in cognitive sciences*, 11(8):333–341, 2007.

James J DiCarlo, Davide Zoccolan, and Nicole C Rust. How does the brain solve visual object recognition? *Neuron*, 73(3):415–434, 2012.

Sander Dieleman, Jeffrey De Fauw, and Koray Kavukcuoglu. Exploiting cyclic symmetry in convolutional neural networks. *arXiv preprint arXiv:1602.02660*, 2016.

Robert Gens and Pedro M Domingos. Deep symmetry networks. In *Advances in neural information processing systems*, pages 2537–2545, 2014.

Ian Goodfellow, Honglak Lee, Quoc V Le, Andrew Saxe, and Andrew Y Ng. Measuring invariances in deep networks. In *Advances in neural information processing systems*, pages 646–654, 2009.

Alexander Kirillov. *An introduction to Lie groups and Lie algebras*, volume 113. Cambridge University Press, 2008.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance and equivalence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 991–999, 2015.

Stéphane Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.

Stéphane Mallat. Understanding deep convolutional networks. *Phil. Trans. R. Soc. A*, 374(2065): 20150203, 2016.

Youssef Mroueh, Stephen Voinea, and Tomaso A Poggio. Learning with group invariant features: A kernel perspective. In *Advances in Neural Information Processing Systems*, pages 1558–1566, 2015.

Michael K Murray and John W Rice. *Differential geometry and statistics*, volume 48. CRC Press, 1993.

Edouard Oyallon and Stéphane Mallat. Deep roto-translation scattering for object classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2865–2873, 2015.

Frédéric Paulin. Groupes et géométrie. *Notes de cours*, 2014.

Jean Petitot. The neurogeometry of pinwheels as a sub-riemannian contact structure. *Journal of Physiology-Paris*, 97(2):265–309, 2003.

Jean Petitot. Neurogéométrie de la vision. *Modeles mathématiques et physiques des architectures fonctionelles. Paris: Éd. École Polytech*, 2008.

Tomaso Poggio, Jim Mutch, Joel Leibo, Lorenzo Rosasco, and Andrea Tacchetti. The computational magic of the ventral stream: sketch of a theory (and why some deep architectures work). 2012.

Anant Raj, Abhishek Kumar, Youssef Mroueh, P Thomas Fletcher, et al. Local group invariant representations via orbit embeddings. *arXiv preprint arXiv:1612.01988*, 2016.

Laurent Sifre and Stphane Mallat. Ecole polytechnique, cmap phd thesis rigid-motion scattering for image classification author. 2014.

Michael Spivak. comprehensive introduction to differential geometry. vol. iv.[a]. 1981.

Christopher Tensmeyer and Tony Martinez. Improving invariance and equivariance properties of convolutional neural networks. *openreview.net/pdf?id=Syfkm6cgx*, 2016.

Thomas Wiatowski and Helmut Bölcskei. A mathematical theory of deep convolutional neural networks for feature extraction. *arXiv preprint arXiv:1512.06293*, 2015.

## Appendix A. Proofs of Theorem 1, Theorem 2 and Lemma

**Theorem 1.**
For all $f \in L^2(\mathbb{R}^2)$, for all $g \in G$,

$$\|\Phi(L_g f) - \Phi(f)\|_2 \leqslant \sqrt{\lambda}\max(1, \sqrt{\|J_g\|}_\infty)\frac{\mu_G((G_0 g)\Delta G_0)}{\mu_G(G_0)}\|f\|_2.$$

*Proof:*

We have

$$\mu_G(G_0)^2\|\Phi(L_g f) - \Phi(f)\|_2^2 = \int_{x\in\mathbb{R}^2}\big(\int_{g'\in G_0} L_{g'g}f(x) - L_{g'}f(x)d\mu_G(g')\big)^2 dx,$$

but

$$\int_{g'\in G_0}(L_{g'g}f(x) - L_{g'}f(x))d\mu_G(g') = \int_{g'\in G_0} L_{g'g}f(x)d\mu_G(g') - \int_{g'\in G_0} L_{g'}f(x)d\mu_G(g'),$$

*i.e.* setting $g'' = g'g$ and using the right-invariance of $\mu_G$,

$$\int_{g'\in G_0}(L_{g'g}f(x) - L_{g'}f(x))d\mu_G(g') = \int_{g''\in G_0 g} L_{g''}f(x)d\mu_G(g'') - \int_{g'\in G_0} L_{g'}f(x)d\mu_G(g').$$

And using $\int_A h - \int_B h = (\int_{A\setminus B} h + \int_{A\cap B} h) - (\int_{B\setminus A} h + \int_{B\cap A} h) = \int_{A\setminus B} h - \int_{B\setminus A} h$, we have

$$\int_{g'\in G_0}(L_{g'g}f(x) - L_{g'}f(x))d\mu_G(g') = \int_{g'\in G_0 g\setminus G_0} L_{g'}f(x)d\mu_G(g') - \int_{g'\in G_0\setminus G_0 g} L_{g'}f(x)d\mu_G(g').$$

Plugging this in the first equation gives

$$\mu_G(G_0)\|\Phi(L_g f) - \Phi(f)\|_2 = \|\int_{g'\in G_0 g\setminus G_0}(L_{g'}f)d\mu_G(g') - \int_{g'\in G_0\setminus G_0 g}(L_{g'}f)d\mu_G(g')\|_2,$$

*i.e.* using a triangle inequality

$$\mu_G(G_0)\|\Phi(L_g f) - \Phi(f)\|_2 \leqslant \|\int_{g'\in G_0 g\setminus G_0}(L_{g'}f)d\mu_G(g')\|_2 + \|\int_{g'\in G_0\setminus G_0 g}(L_{g'}f)d\mu_G(g')\|_2.$$

Now observe that by interverting the integrals using Fubini's theorem,

$$\|\int_{g'\in G_0\setminus G_0 g}(L_{g'}f)d\mu_G(g')\|_2 = \sqrt{\int_{g_1\in G_0\setminus G_0 g}\int_{g_2\in G_0\setminus G_0 g}\big(\int_{x\in\mathbb{R}^2}(L_{g_1}f)(x)(L_{g_2}f)(x)dx\big)d\mu_G(g_1)d\mu_G(g_2)},$$

and using a Cauchy-Schwarz inequality,

$$\|\int_{g'\in G_0\setminus G_0 g}(L_{g'}f)d\mu_G(g')\|_2 \leqslant \sqrt{\int_{g_1\in G_0\setminus G_0 g}\int_{g_2\in G_0\setminus G_0 g}\|L_{g_1}f\|_2\|L_{g_2}f\|_2 d\mu_G(g_1)d\mu_G(g_2)}.$$

As for all $g' \in G_0$ we have $\|L_{g'} f\|_2 = \|f\sqrt{|J_{g'}|}\|_2 \leqslant \sqrt{\lambda}\|f\|_2$ with a change of variables, we have

$$\| \int_{g' \in G_0 \backslash G_0 g} (L_{g'} f) d\mu_G(g') \|_2 \leqslant \sqrt{\lambda} \mu_G(G_0 \setminus (G_0 g)) \|f\|_2.$$

For the other term, note that by setting $g'' = g'g^{-1}$, we have

$$\| \int_{g' \in G_0 g \backslash G_0} (L_{g'} f) d\mu_G(g') \|_2 = \| \int_{g'' \in G_0 \backslash G_0 g^{-1}} (L_{g''g} f) d\mu_G(g'') \|_2 = \| \int_{g'' \in G_0 \backslash G_0 g^{-1}} (L_{g''} L_g f) d\mu_G(g'') \|_2,$$

and then similarly,

$$\| \int_{g'' \in G_0 \backslash G_0 g^{-1}} (L_{g''} L_g f) d\mu_G(g'') \|_2 = \sqrt{\int_{g_1 \in G_0 \backslash G_0 g^{-1}} \int_{g_2 \in G_0 \backslash G_0 g^{-1}} \|L_{g_1} L_g f\|_2 \|L_{g_2} L_g f\|_2 d\mu_G(g_1) d\mu_G(g_2)}.$$

As for all $g' \in G_0$ we have $\|L_{g'} L_g f\|_2 = \|f\sqrt{|J_{g'g}|}\|_2 \leqslant \sqrt{\lambda \|J_g\|}_\infty \|f\|_2$, we have

$$\| \int_{g' \in G_0 g \backslash G_0} (L_{g'} f) d\mu_G(g') \|_2 \leqslant \sqrt{\lambda \|J_g\|}_\infty \mu_G(G_0 \setminus (G_0 g^{-1})) \|f\|_2.$$

Therefore

$$\mu_G(G_0) \|\Phi(L_g f) - \Phi(f)\|_2 \leqslant \sqrt{\lambda \|J_g\|}_\infty \mu_G(G_0 \setminus (G_0 g^{-1})) \|f\|_2 + \sqrt{\lambda} \mu_G(G_0 \setminus (G_0 g)) \|f\|_2,$$

and the following fact concludes the proof:

$$\mu_G(G_0 \setminus (G_0 g^{-1})) + \mu_G(G_0 \setminus (G_0 g)) = \mu_G((G_0 g) \setminus G_0) + \mu_G(G_0 \setminus (G_0 g)) = \mu_G((G_0 g) \Delta G_0).$$

□

**Theorem 2.**
For all $f \in \mathcal{X}$, for all $\xi, \xi' \in \mathfrak{g}$,

$$\|[\tilde{X}_\xi, \tilde{X}_{\xi'}]_{\Phi(f)}\|_2^2 \leqslant \lambda \big[ \frac{d}{ds}_{|s=0} \frac{\mu_G((G_0 \exp(s[\xi, \xi'])) \Delta G_0)}{\mu_G(G_0)} \big]^2 \|f\|_2^2.$$

*Proof:*

As $\Phi$ realizes a diffeomorphism from $G \cdot f$ onto its image, and as $\Phi$ equals its differential from *Lemma*, we have that for all vector field $X$ on $G \cdot f$, $\Phi_*(X)(\Phi(f)) = (d\Phi)_f(X(f)) = \Phi(X(f))$. Hence

$$\begin{aligned}
[\tilde{X}_\xi, \tilde{X}_{\xi'}]_{\Phi(f)} &= [\Phi(X_\xi), \Phi(X_{\xi'})]_f \\
&= [\Phi(X_\xi) \circ \Phi^{-1}, \Phi(X_{\xi'}) \circ \Phi^{-1}]_{\Phi(f)} \\
&= [\Phi_*(X_\xi), \Phi_*(X_{\xi'})]_{\Phi(f)} \\
&= \Phi_*([X_\xi, X_{\xi'}])(\Phi(f)) \\
&= \Phi([X_\xi, X_{\xi'}]_f).
\end{aligned}$$

Recall that the Lie bracket of left-invariant vector fields is given by the opposite of the Lie bracket of their corresponding generators, hence in our case:

$$[X_\xi, X_{\xi'}] = X_{-[\xi,\xi']} = -X_{[\xi,\xi']}.$$

Therefore,

$$\begin{aligned}
\|[\tilde{X}_\xi, \tilde{X}_{\xi'}]_{\Phi(f)}\|_2 &= \|\Phi([X_\xi, X_{\xi'}]_f)\|_2 \\
&= \|\Phi(X_{[\xi,\xi']}(f))\|_2 \\
&= \|\Phi(\lim_{t\to 0} \frac{1}{t}(L_{\exp(t[\xi,\xi'])}f - f))\|_2 \\
&= \|\lim_{t\to 0} \frac{1}{t}\big(\Phi(L_{\exp(t[\xi,\xi'])}f) - \Phi(f)\big)\|_2.
\end{aligned}$$

From *Theorem 1*, we have

$$\|\Phi(L_{\exp(t[\xi,\xi'])}f) - \Phi(f)\|_2 \leqslant \sqrt{\lambda}\max(1, \sqrt{\|J_{\exp(t[\xi,\xi'])}\|}_\infty)\frac{\mu_G((G_0\exp(t[\xi,\xi']))\Delta G_0)}{\mu_G(G_0)}\|f\|_2.$$

As $\exp(t[\xi,\xi']) \to e$ when $t \to 0$, its Jacobian goes to 1. Moreover, as $f$ has a gradient with fast decay, we can take the limit out of the $L^2$-norm, which concludes the proof.

□

**Lemma.**
For all $f \in \mathcal{X}$ and $\xi \in \mathfrak{g}$,

$$\frac{d}{dt}_{|t=0}\Phi(L_{\exp(t\xi)}f) = \Phi\big(\frac{d}{dt}_{|t=0}(L_{\exp(t\xi)}f)\big).$$

*Proof:*

For all $x \in \mathbb{R}^2$,

$$\begin{aligned}
\big(\frac{d}{dt}_{|t=0}\Phi(L_{\exp(t\xi)}f)\big)(x) &= \frac{d}{dt}\big(\frac{1}{\mu_G(G_0)}\int_{g'\in G_0}(L_{g'}L_{\exp(t\xi)}f)(x)d\mu_{G_0}(g')\big)_{|t=0} \\
&= \frac{1}{\mu_G(G_0)}\int_{g'\in G_0}\frac{d}{dt}\big(f(\exp(-t\xi)g'^{-1}x)\big)_{|t=0}d\mu_{G_0}(g') \\
&= \frac{1}{\mu_G(G_0)}\int_{g'\in G_0} d_{(g'^{-1}x)}f\big(-\xi(g'^{-1}x)\big)d\mu_{G_0}(g') \\
&= \Phi\big(d.f\big(-\xi(\cdot)\big)\big)(x) \\
&= \Phi\big(\frac{d}{dt}_{|t=0}(L_{\exp(t\xi)}f)\big)(x).
\end{aligned}$$

□

## Appendix B. Supplementary material

The next three propositions are taken from the publicly available french textbook Paulin (2014), in which they're respectively numbered as E.7, 1.60, 1.62.

**Proposition B.0** Let $G$ be a Lie group and $\rho : G \to GL(V)$ a finite-dimensional Lie group representation of $G$. Then for all $v \in V$, the map defined by $g \in G \mapsto \rho(g)v$ has constant rank, and the stabilizer $G_v$ is an embedded Lie subgroup of $G$.

**Proposition B.1**. Let $G$ be a Lie group, $H$ be an embedded Lie subgroup of $G$, and $\pi : G \to G/H$ be the canonical projection. There exists one and only one smooth manifold structure on the topological quotient space $G/H$ turning $\pi$ into a smooth submersion. Moreover, the action of $G$ on $G/H$ is smooth, and if $H$ is normal in $G$, then $G/H$ is a Lie group, $\pi$ is a Lie group morphism, the Lie algebra $\mathfrak{h}$ of $H$ is an ideal of the Lie algebra $\mathfrak{g}$ of $G$ and the linear map from $T_eG/T_eH$ to $T_{eH}(G/H)$ induced by $T_e\pi$ is a Lie algebra isomorphism from $\mathfrak{g}/\mathfrak{h}$ to the Lie algebra of $G/H$.

**Proposition B.2**. Let $M$ be a manifold together with a smooth action of a Lie group $G$, and $x \in M$; *(i)* the canonical map $\Theta_x : G/G_x \to M$ defined by $\Theta_x(gG_x) = gx$ is a one-to-one immersion, whose image is the orbit $G \cdot x$; *(ii)* the orbit $G \cdot x$ is a submanifold of $M$ if and only if it is locally closed in $M$; *(iii)* if $G \cdot x$ is locally closed, then $\Theta_x$ is a diffeomorphism from $G/G_x$ to $G \cdot x$.